

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☐ **BLACK BORDERS**
- ☐ **IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- ☐ **FADED TEXT OR DRAWING**
- ☐ **BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- ☐ **SKEWED/SLANTED IMAGES**
- ☐ **COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- ☐ **GRAY SCALE DOCUMENTS**
- ☐ **LINES OR MARKS ON ORIGINAL DOCUMENT**
- ☐ **REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- ☐ **OTHER:** _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : H04N 5/765	A1	(11) International Publication Number: WO 98/10586 (43) International Publication Date: 12 March 1998 (12.03.98)
(21) International Application Number: PCT/GB97/02383 (22) International Filing Date: 5 September 1997 (05.09.97) (30) Priority Data: 9618556.6 5 September 1996 (05.09.96) GB (71) Applicant (for all designated States except US): DISCREET LOGIC INC. [CA/CA]; 10 Duke Street, Montreal, Quebec H3C 2L7 (CA). (72) Inventor; and (75) Inventor/Applicant (for US only): TYSON, Peter, John [GB/GB]; Discreet Logic, Mill House, Overbridge Square, Hambridge Lane, Newbury RG14 5UX (GB). (74) Agent: ATKINSON, Ralph; Atkinson & Co., The Technology Park, 60 Shirland Lane, Sheffield S9 3SP (GB).		(81) Designated States: CA, JP, US, European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE). Published <i>With international search report.</i> <i>Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i>

(54) Title: IMAGE PROCESSING SYSTEM CONVERTING BETWEEN VIDEO DATA AND ADDRESSED DATA

(57) Abstract

Video circuits (204) are arranged to convey sequential video signals, usually in the form of interlaced fields with each field being separated by a field blanking period. Data circuits (203) are arranged to convey randomly addressable image data and do not include blanking. A processing system (202) is arranged to convert between video data and address data to effect transfers between said video circuits and said data circuits at a rate equal to or greater than video rate. Data may be read from a digital video tape recorder and written to local storage in a format compatible with the addressed environment continuously at video rate.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

IMAGE PROCESSING SYSTEM CONVERTING BETWEEN VIDEO DATA AND ADDRESSED DATA**Technical Field**

The present invention relates to image data processing, wherein video
5 circuits are arranged to convey sequential video signals and data circuits are arranged to convey randomly addressable image data.

Background Art

Systems for image data processing are known in which video data is
10 transmitted as a digital stream clocked with reference to synchronization signals representing the start of frames or the start of interlaced fields within said frames. In broadcast environments and within facilities houses, the D1 standard for the transmission of digital video is popular, wherein the video signal is transmitted as a luminance signal (Y) plus color difference signals
15 (UV) to facilitate conversion to conventional broadcast standards such as NTSC and PAL. The digital signals are clocked in real-time at video rate to facilitate conversion in real-time to analog signals for transmission purposes. Thus, within such facilities, the transmission of digitized video signals is similar to the transmission of their analog equivalents, where the rate of
20 transmission remains constant and all transmissions are synchronized to the field and frame blanking intervals.

Image data processing systems are also known in which image data is transmitted in a way substantially similar to the transmission of any other type of data within a computer environment. Within computer environments, it
25 is unusual for broadcast quality video signals to be produced, therefore, conventionally, signals have been transmitted as red, green and blue (RGB) color components so as to be compatible with CRT monitoring devices. Data signals within a computing environment are also transmitted relatively asynchronously and relatively randomly, given that each data transfer usually
30 takes place by issuing address signals, wherein said address signals may represent locations within solid state memory devices or locations within

mass storage devices such as magnetic disks etc. Similarly, asynchronous protocols have been developed for the transmission of signals between such devices, such as the small computer systems interface (SCSI).

In systems developed for the transmission and processing of YUV D1
5 signals, reliance has been made significantly upon purpose built hardware. However, post-production and broadcast facilities are tending to move towards software intensive solutions on relatively low cost general purpose workstations. This allows software applications (defining a system's
10 functionality) to be developed independently of hardware platforms, facilitating modification and upgrading when compared to hard-wired hardware solutions. However, although workstations are becoming more powerful, many relatively simple tasks may place unprecedented burdens upon central processing units and a high performance software package may become significantly downgraded by processing time being occupied by
15 relatively low-level housekeeping procedures.

Summary of the Invention

According to a first aspect of the present invention, there is provided image data processing apparatus for converting between sequential video
20 data and addressed image data to effect transfers between video circuits and data circuits, said apparatus comprising processing means arranged to transfer data between said circuits at a rate equal to or greater than video rate; and formatting means operable to rearrange the order of components of said data representing color samples being transferred by said processing
25 means.

Preferably, said formatting means includes packing means, arranged to pack color components together into data words so as to remove redundant regions of said words.

According to a second aspect of the present invention, there is
30 provided a method of processing image data for converting between sequential video data and addressed image data to effect transfers between

video circuits and data circuits, said method comprising steps of transferring data between said circuits at a rate equal to or greater than video rate; and rearranging the order of components of said data representing color samples being transferred.

- 5 Preferably, a color-space conversion process converts data between luminance plus color difference representations and primary color representations.

 According to a third aspect of the present invention, there is provided image data processing apparatus for converting between sequential video
10 data and addressed image data to effect transfers between video circuits and data circuits, said apparatus comprising processing means arranged to transfer data between said circuits at a rate equal to or greater than video rate; and buffering means arranged to buffer transfers between said video circuits and said addressed circuits, said buffering means including a first
15 buffering unit to provide access to a transmission environment and a second buffering unit to provide access to a storage environment.

Brief Description of the Drawings

 Figure 1 shows a post-production video artist using an application for
20 modifying image frames, including an applications platform, a video tape recorder and an image processing system providing real-time communication between the applications platform and the tape recorder;

 Figure 2 identifies a schematic representation of the environment illustrated in Figure 1, in which an image processing system converts
25 between sequential line-based video signals and randomly addressable image data;

 Figure 3 illustrates differences between line-based video signals and address-based image data;

 Figure 4 details the image processing system shown in Figure 2,
30 including a video buffer, a router, a color space converter, a proxy generator, a re-formatter, a disc buffer, a network buffer, a parity generator and a PCI

bus;

Figure 5 details the video buffer identified in Figure 4;

Figure 6 details the router identified in Figure 4;

Figure 7 illustrates the configuration of PCI devices, including the PCI
5 bridges shown in Figure 4;

Figure 8 details the color-space converter shown in Figure 4;

Figure 9 details the proxy generator identified in Figure 4;

Figure 10 details the re-formatting circuit identified in Figure 4,
including a packing circuit; and

10 Figure 11 details the packing circuit identified in Figure 10.

Detailed Description of the Preferred Embodiment

The invention will now be described by way of example only with
reference to the previously identified drawings. A post-production facility is
15 illustrated in Figure 1, in which a video artist 101 is seated at a processing
station 102. Images are displayed to the artist via a visual display unit 103
and manual selections and modifications to the displayed images are
effected in response to manual operation of a stylus 104 upon a touch tablet
105. In addition, a conventional keyboard 106 is provided to allow
20 alphanumeric values to be entered directly. The monitor 103, tablet 105 and
keyboard 106 are interfaced to an image manipulating workstation 107, such
as an Indigo Max Impact, manufactured by Silicon Graphics Inc., running
compositing applications, such as "FLINT" or "FLINT RT" licensed by the
present applicant.

25 Image data may be supplied to the workstation 107 from a D1 digital
video tape recorder 108 via an image processing system 109. The video tape
recorder 108 and the processing system 109 are both controlled directly in
response to commands issued by the artist 101, thereby effectively
embedding the operation of these machines within the applications
30 environment. Processing system 109 is arranged to receive video data from
the video recorder 108 at video rate and is arranged to write said data to its

own internal storage devices at this rate. The processing system 109 is then in a position to make this recorded data available to the workstation 107, or to similar devices via a high bandwidth network such as "HiPPI", via a network cable 110.

5 The environment shown in Figure 1 is illustrated schematically in Figure 2. The workstation 107, its interfaces and its associated applications may be considered as a local applications platform 201. The processing system 109 may be considered as comprising an image processing system 202 having associated circuitry which may be considered as belonging to an
10 addressed data environment 203 or to a video environment 204. The local applications platform 201 communicates with the image processing system 202 via the addressed environment 203. The addressed environment 203 also communicates with a local array of discs 205, which may be configured in accordance with RAID protocols. Thus, the local array 205 may include a
15 conventional SSA adapter, such as the type supplied by Pathlight Technology Inc. of 767 Warren Road, Ithaca, New York, 14850 and the addressed environment 203 includes circuitry for transmitting and receiving data from the SSA adapter in accordance with conventional protocols. The addressed environment 203 also includes interface cards for connection to a
20 HiPPI network 206.

A D1 serial digital input 207 supplies synchronised D1 video to the video environment 204. Similarly, output video from the video environment 204 is supplied to an output cable 208. Interfaces 207 and 208 may be connected to a video tape recorder, such as tape recorder 108 shown in
25 Figure 1. The video environment 204 also supplies analog video signals to a video monitor 209, allowing an operator to view video images as they are being transmitted through the image processing system.

The addressed environment 203 and the video environment 204 have two distinct differences in terms of the way in which data is processed within
30 these environments. Within the video environment 204, pixel data is transmitted sequentially starting from the top left position of an image frame

and then scanning from left to right along a scan line, with lines being transmitted sequentially from top to bottom. Thus, the location of a particular pixel within an image frame is determined by its temporal location within a transmitted frame's worth of data. In this way, it is not necessary to provide
5 addressing channels given that the order of the pixels within a frame always follows the same predetermined scanning pattern, usually consisting of interlaced frames. However, within the addressed environment 203 data need not be transmitted in a sequential way. Each data transfer is accomplished by identifying its address within storage space, therefore a
10 complete frame's worth of data may be transmitted in any order provided that each pixel data is accompanied by addressing information. This mechanism for transferring data is particularly advantageous when using disk arrays which, due to mechanical differences between the drives, will result in data being transmitted in a substantially random way. However, provided that
15 each unit of data is accompanied by its corresponding address information, simple buffering techniques allow the full frame to be re-established in its original pixel order.

A second significant difference between the addressed environment 203 and the video environment 204 is that the transmission of data within the
20 video environment is synchronized such that data is transmitted at 13.5 megahertz or integer multiples thereof. Thus, data may be transmitted within the video environment 204 at 27 megahertz such that each channel is capable of conveying two real-time video streams. In this way, after a predetermined period of time, the amount of data transferred will be known
25 and will remain fixed. This differs from the addressed environment 203 in which, although provided with a notional transfer bandwidth, such as 33 megahertz in a typical PCI system, the actual availability of buses for a specific data transfer will depend upon other constraints placed upon the environment. Thus, if the bus is required to effect other transfers within the
30 environment or if a bus mastering processor is tied up with other activities, transmission will cease until the process can again obtain access to the bus.

Thus, transfers are not synchronized to specific events but are interrupt driven, with priority given to the highest priority interrupt.

These two differences create problems when attempting to effect real-time data transfers between a video environment and a typical addressed environment. In known systems, transfers often take place at less than real-time, with a block of data being transferred in the first protocol, whereafter processing routines are required to convert the data into a different protocol. Under such circumstances, the video environment must be modified and interrupted so as to allow it to transfer data to the addressed environment. In the arrangement shown in Figure 2, the video environment is arranged to transfer data transparently to an environment it sees as being video compatible. Similarly, the addressed environment 203 may transfer data with the image processing system 202 in a way which it perceives as being addressed-compatible. These transfers occur at a rate which is at least equal to video rate and possible higher, such that the operations of a transmitting device are not limited by the capabilities of a receiving device. Consequently, the transfer of data from a D1 digital video environment to an RGB addressed environment occurs with no more difficulty than a similar transfer from a D1 digital environment to a similarly configured processing environment, arranged to process D1 digital signals directly.

Differences between the transmission of digital video and addressed data are illustrated in Figure 3. An analog video picture is illustrated at 301 and is made up of a first field of scan lines 302, shown as solid lines, which are interlaced with a second field of scan lines, illustrated by broken lines 303. A first field scan consists of the first field of lines 302 being scanned from the top left of the picture to the bottom right of the picture. A field blanking interval then occurs as electronics are reset to facilitate the operation of the next scan. The subsequent field contains the interlaced lines 303 and again a scanning operation is performed from the top left to the bottom right of the picture. The scanning of two interlaced fields in this way constitutes a whole frame and after a frame of two fields have been scanned,

the process is repeated at a frame rate of 25 or 30 frames per second.

D1 serial digital signals may be generated in real-time by receiving a scanning analog signal which is supplied to an analog to digital converter 304. The analog video signal is conveyed as a luminance signal Y plus two color difference signals U and V and each of these samples are digitized and multiplexed to produce a serial stream at 13.5 megahertz. These serial digits are illustrated at 305 and represent a stream of data corresponding to the luminance and color of the original analog signal.

The digital stream 305 is synchronously clocked in response to clocking signals 306, generated in synchronism with the field blanking intervals. Thus, although the video information, originally in analog form, has been digitized into a data stream, this data stream still synchronously follows the field intervals of the original source video. Similarly, these digital samples represent color components which are similar to those used in analog video environments, defining luminance plus color difference components. In this way, real-time conversion is facilitated between digital and analog environments with the video signal having many characteristics in common with conventional analog transmission.

An analog picture display within a computing environment is illustrated at 308. A computer CRT receives analog signals for its red, green and blue components, which are generated in response to digital input signals. The storage of data within such environments is therefore carried out entirely within RGB color space, with no requirement being present for conversion into YUV color space. Consequently, when receiving conventional video data, although in digitized form, it is necessary to perform this conversion from YUV color space to RGB color space.

A visual display is generated by scanning a frame buffer 309 at an appropriate frame display rate. A computer system may refresh its screen at whatever rate is considered appropriate and, in order to enhance the quality of CRT displays, computer images tend to be refreshed at a higher rate than that of television broadcasts, with frames tending to be non-interlaced. Thus,

a computer image tends not to consist of non-interlaced frames and the refresh rate tends to be higher than that occurring in conventional broadcast television systems.

In conventional television and video systems, new data is received for each displayed frame and the system is described as operating in real-time or at video rate. In computer systems this is often not the case and although a screen may be refreshed at a higher rate than video systems, the rate at which new data is used may be less than that occurring within video systems. Thus, a computer system has the option to display the same source data many times by reading the data from it's frame buffer 309.

Frame buffers are known within video environments and usually involve the writing of data to the buffer in sequential order with the data then being read from the buffer in sequential order, unless the buffer is being used to perform a video effect. In computer based systems, the transfer of data does not necessarily take place in a sequential way. In the system shown in Figure 3, the data is read sequentially from frame buffer 309 in order to generate display 308. However, data may be written to the frame buffer 309 in any random order and at a rate which may depend upon the processing capabilities of the central processing unit itself. Thus, data may be written to frame buffer 309 by placing an address on address bus 310 so as to identify a particular location within the frame buffer. Thereafter data is supplied to data bus 311 so as to allow the data to be written to the location defined by the address bus. In this way, it is possible to write data to the frame buffer 309 in a completely random way because each item of data is identified by it's own specific address supplied to the address bus 310.

The rate at which data is clocked on address and data buses 310, 311 will depend upon the system employed but essentially it may be considered as being substantially asynchronous compared to the clocking of video data through, say, analog to digital converter 304. Furthermore, in more powerful machines, the data rate at which data is clocked through the internal buses will increase with system power and complexity. Thus, within the data

environment, the rate of transfer will tend to depend upon the system's capabilities and data will be transmitted at a maximum possible rate. This compares with the video environment in which the rate of transfer is dependant upon the nature of the data itself, with data usually being
5 transferred at conventional video rate or, in more sophisticated systems, at integer-multiples thereof.

Image processing system 202 is detailed in Figure 4. The processing system includes a video buffer 401 which communicates with a synchronously clocked router 402. The router 402 is arranged to direct data
10 to a color-space converter 403, a proxy generator 404 and a digital to analog converter 405, supplying analog RGB signals to video monitor 209. The router 402 also communicates with reformatting circuits 406 and 407, which in turn communicate with their respective buffers 408 and 409. Disk buffer 408 facilitates communication between the router 402 and local disks 205;
15 and as such it may be identified as the disk buffer. The local RAID 205 includes redundant parity information, therefore parity circuit 410 is associated with disk buffer 408 in order to generate parity as data is being written to the disk array 205 and, where necessary, to reconstitute lost data from said parity information. Parity circuit 410 and disk buffer 408
20 communicate with the local RAID 205 via an SSA interface card 411.

Communications with network 206 are effected via a HiPPI network card 412, while communications with the local applications platform 201 takes place via a SCSI interface card 413, preferably provided with two SCSI channels in the preferred embodiment.

25 The processing system 202 is controlled by a programmable processing unit 414, which is responsible for co-ordinating activities within the processing system and for downloading instructions to specific components within the processing system. In the preferred embodiment, processing unit 414 is implemented as an Intel microprocessor communicating with a primary
30 thirty-two bit PCI bus 415 clocked at 33 megahertz. The primary PCI bus 415 allows processing unit 414 to supply memory mapped control signals to

associated processing sub-systems. However, the primary bus 415 is not used for the transmission of image data. The mechanism for transferring image data within the image processing system 202 is the router 402, with transfers also taking place via one or both of buffers 408 and 409.

5 Network buffer 409, network card 412 and interface card 413 communicate via a secondary PCI bus 416, which may be considered as a network bus. Secondary bus 416 is connected to primary bus 415 via a PCI bridge 417. Bridge 417 is configured to allow control information to be transmitted from primary bus 415 to secondary bus 416 as if the bridge 417
10 effectively did not exist. However, data lying outside a specified address range will be treated as data and as such bridge 417 will be perceived as being closed. Consequently, any image data supplied to secondary bus 416 can communicate between network card 412, interface card 413 and network buffer 409 but cannot be conveyed to the primary bus 415 via the bridge 417,
15 which will be seen as open.

A similar arrangement is provided for communication between the disk buffer 408 and the disk interface 411. A secondary PCI bus, which may be considered as the disk bus 418 is connected to the primary PCI bus 415 via a second PCI bridge 419. Bridge 419 allows control information to be
20 transferred from the processing unit 414 to the interface card 411, its associated SSA adapter and to disk buffer 408. However, the bridge 419 is effectively closed for the transmission of image data, such that image data supplied to the network bus 418 is blocked from reaching the primary bus 415. Consequently, no major burdens are placed upon the processing unit
25 414 and its associated primary bus. Processing unit 414 is only concerned with configuring other subsystems and is not itself directly responsible for controlling the transfer of image data via bus mastering or other techniques.

Data transfers within the image processing system 202 preferably take place within RGB color space. Consequently, D1 video signals supplied to
30 the video environment 204 are color-space converted within said environment using conventional dedicated circuits employed in digital video

systems. However, signals supplied to the processing system 202 from the video environment are sequentially clocked, consist of interlaced fields and include field blanking and would normally be perceived as video signals. The addressed environment 203 includes the SSA adapter for supplying data to the local raid 205. Data supplied to the raid 205 is effectively data like any other type of data and, as such, the fact that it represents image frames is immaterial to the operation of the SSA environment.

Typical operations of the processing system shown in Figure 4 may be considered as follows. In a first mode, a local application running on platform 201 may request a search and preview of information stored on a video tape recorder, such as recorder 108. The video information will be read from the video tape recorder and supplied to the video buffer 401. The video buffer 401 is responsible for converting the serial data into a parallel stream and as such may be considered as a serial to digital interface. In addition, the video buffer 401 removes field blanking such that the nature of the data supplied to the router 402 from the video buffer 401 ceases to be "video" in the accepted sense. The router 402 routes the incoming signal to digital to analog converter 405, allowing the incoming video clip to be viewed on monitor 209. In parallel to this, the incoming video stream is supplied to proxy generator 404 configured to produce a reduced bandwidth version of the incoming video (with a resolution reduced by half in both dimensions) which is in turn routed to the network buffer 409 where it is written sequentially to said buffer so as to build up a complete frame of a reduced bandwidth image.

Network buffer 409 is configured as a double frame buffer therefore as one frame buffer receives data from router 402, the other frame buffer is available to supply image data to the SCSI interface card 403 via the network bus 416.

After a particular clip has been identified by the above search procedure, the clip itself may now be written to the local RAID 205. In parallel with the procedures identified above, resulting in the incoming clip being

monitored on video monitor 209 and also previewed in proxy form on display 103, router 402 also routes the full bandwidth version of the incoming video to the disk buffer 408. Disk buffer 408 includes two frame buffers, to effect double buffering at the frame level, allowing the first frames-worth of data to
5 be written to the buffer while a second frames worth of data is read from said buffer. Data in disk buffer 408 is striped and an additional parity stripe is generated by parity circuit 410, whereafter all of the stripe data is written to the RAID 205 via the SSA adapter and interface card 411.

A third mode involves clip output, where image data is read from the
10 local RAID 205 and supplied to a video store, such as a video tape recorder. From the disks and the SSA adapter, data is written to the disk buffer 408. This transfer will take place randomly, due to the way in which the SSA adapter will package data received from the array. The data is therefore assembled within the disk buffer 408 and read from said buffer sequentially to
15 provide a stream to the router 402. The clip will have been stored in RGB format therefore it is necessary to direct the signal to the color-space converter 403, arranged to reconvert said signal back to YUV format for application to the video buffer 401. In parallel with this, the RGB signal is supplied to the video monitor via converter 405 and the transfer is viewed on
20 display 103 by directing the RGB signal to the proxy generator 404, again via router 402, whereafter said reduced band-width signal is supplied to the network buffer 409.

A fourth mode consists of clip playback; substantially similar to clip output but without the need to direct a color-space converted signal to the
25 video buffer 401.

As an alternative to supplying reduced bandwidth signals to interface card 413, router 402 may be arranged to direct full bandwidth RGB signals to network buffer 409, whereafter image data read from said buffer is directed to the network card 412 for application to the high band-width network.

30 Unlike the video environment, the addressed environment is equally capable of transmitting individual frames in addition to transmitting video

clips. This involves a transfer between interface card 411 and interface cards 412 or 413. As previously stated, a transfer of this type does not take place via the primary PCI bus 415. During a frame read, individual frame data from interface card 411 is written to disk buffer 408, read from said buffer, applied
5 to router 402, routed to network buffer 409 and read from said buffer 409. If the frames are to be supplied to interface card 413 at reduced bandwidth, the signal is routed via the proxy generator 404 and then routed to the network buffer 409.

Similarly with a frame write, an incoming frame may be received by
10 interface card 413 and written to network buffer 409. The frame is then read from network buffer 409 and written to disk buffer 408 via the router 402. The frame data is then read from disk buffer 408, parity is generated and the resulting stripes are supplied to interface card 411.

Video buffer 401 effectively consists of two buffers each arranged to
15 convey two real-time video streams to router 402 at 27 megahertz. Each of these individual buffer circuits may therefore simultaneously receive a D1 video stream at 13.5 megahertz while transmitting a similar stream at 13.5 megahertz.

The video buffer 401 is detailed in Figure 5, consisting of a first buffer
20 circuit 501 and a substantially similar second buffer circuit 502. The first buffer circuit 501 will be described and it should be understood that substantially similar operations are effected by the second buffer circuit 502.

An incoming D1 video stream, color converted to RGB, is received on an input serial line 503. The incoming data may include an audio embedded
25 stream and includes field blanking. The audio inbedded stream is separated by audio embedding circuit 504 and supplied to an audio buffer 505. A switch 506 directs incoming video data to a first field store 507 or to a second field store 508. Each field store is arranged to store only video data and does not store field blanking. Thus, the process of writing the serial stream to one of
30 said stores effectively removes the field blanking from the video stream such that, thereafter, the data is transmitted as substantially contiguous blocks.

The field buffers 507 and 508 provide double buffering such that as data is written to the second field buffer 508, in accordance with the configuration shown in Figure 5, data previously written to the first field buffer 507 may be read as parallel thirty-two bit words at twenty-seven megahertz
5 for application to the router 402 over bus 509. The reading process will also access audio buffer 505, thereby adding audio data to the twenty-seven megahertz data stream.

Within a field period, it is also possible for data to be received from bus 509 for application to output serial digital link 510. The field period is divided
10 into two sub-periods, within the twenty-seven megahertz domain, and in said second sub-period audio data may be written to audio buffer 511, with a field of video data being written to field store 512 or field store 513. Under the configuration shown in Figure 5, incoming data is written to the second field store 513, allowing the first field store 512 to be read serially at 13.5
15 megahertz to provide a serial stream to the audio embedding circuit 514. At circuit 514 audio data is embedded in accordance with the AES protocol by reading audio data from audio buffer 511. Interlaced RGB video with field blanking at 13.5 megahertz is then supplied to output channel 510. Thus, the reading of field buffers 512 or 513 is appropriately delayed in order to
20 introduce the required field blanking intervals.

The router 402 is detailed in Figure 6 and is fabricated around six thirty-two bit buses clocked at twenty-seven megahertz. The transfer of image data in this mode represents the preferred transmission protocol within the processing system. It is conveyed along the parallel bus, similar to data
25 transmission but this bus is synchronised at twenty-seven megahertz and does not require an associated address bus. A first thirty-two bit bus 601 receives networked data from the reformatting device 407. The second thirty-two bit bus 602 receives disk information from the storage devices via reformatting circuit 406. The third bus 603 receives a first video stream from
30 video buffer 401, while the thirty-two bit bus 604 receives the second video stream from data buffer 401. The fifth thirty-two bit bus 605 receives the

output from the color-space converter 403, with the sixth bus 606 receiving a similar output from the proxy generator 404.

Routing is affected via the buses because in addition to the six input sources, seven output destinations are connected to the bus. The first
5 selector 607 receives input from the disk bus 602, the first video bus 603, the second video bus 604, and the proxy bus 606. Selector 607 receives instructions from the processing unit 414 to select one of these sources thereafter the selected source is applied to the network reformatting circuit 407.

10 A second selector 608 receives input from the network bus 601, the first video bus 603, the second video bus 604 and the proxy bus 606. Again, in response to control signals from the processing unit 414, selector 608 is arranged to select one of these input signals by application to the disk reformatting circuit 406.

15 Communication paths between the router 402 and the video buffer 401 are bi-directional and are configured so as to transmit two real-time video sources over a twenty-seven megahertz transmission channel. To achieve this, one of the sources will be supplied to the router with the second multiplexed signal being supplied from the router back to the video buffer
20 401. The router therefore includes a first multiplexor 614 and a second multiplexor 615 each arranged to connect multiplexed channels to respective input or output ports within the router. A third selector 609 receives inputs from the network bus 601, the disk bus 602, color space converter bus 605 and the proxy bus 606. A selection is made by selector 609, in response to
25 control instructions from the processing unit 414, resulting in a selected input signal being supplied to the multiplexor 614. Similarly, fourth selector 610 receives inputs from the network bus 601, the disk bus 602, the color space converter bus 605 and the proxy bus 606. Again, in response to control signals issued by the processing unit 414, a selected signal is supplied to
30 multiplexor 615.

A fifth selector receives inputs from the network bus 601 and the disk

bus 602. Again, control signals are received from the processing unit 414 so as to select one of these input signals which is in turn supplied to the color-space converter 403.

Inputs from the first video bus 603, the second video bus 604 and the
5 proxy bus 606 are supplied to a sixth selector 612. In response to control signals from the processing unit 414, the sixth selector 612 supplies a selected signal to the proxy generator 414. The seventh selector 613 receives inputs from the first video bus 603 and the second video bus 604. An output is selected in response to control signals from the processing unit
10 414, resulting in the selected signal being supplied to the digital to analog converter 405.

It can be appreciated that the router achieves a routing function by allowing a signal to be selected from its transmission bus. In this way, a switch is effectively non-blocking because the transmission of one signal
15 along its respective bus cannot effect the transmission of other signals along their respective buses. The router does not provide for all possible interconnections and is tailored to meet the requirements of the system's overall functionality. However, additional routing paths may be introduced by allowing signals to bypass through the proxy generator and/or the color-
20 space converter.

Data is transmitted to interface cards 412, 413 and 411 in accordance with PCI protocols. The PCI environment consists of a primary PCI bus 415 with secondary PCI buses 416 and 418 connected to said primary bus by respective PCI bridges 417 and 419. The processing unit 414 provides the
25 primary bus master for the PCI environment, although other devices, such as the SSA adapter associated with the disk drives, may be allowed to bus master in preference to this processing unit. When operating power is initially supplied to processing unit 414, configuration instructions are automatically retrieved from associated read-only memory and these instructions will
30 determine which PCI devices are connected to the primary bus, along with an identification of their configuration requirements. This process is known in

the art as scanning or probing the bus and in order to facilitate this process PCI devices implement a base set of configuration registers, in additions to device-specific configuration registers.

5 The configuration instructions read a sub-set of the devices
configuration registers in order to determine the presence of the device and
it's type. Having determined the presence of the device, other configuration
registers for the device are accessed to determine how many blocks of
memory and the degree of input/output space is required in order to effect
satisfactory operation. Memory associated with the device is then
10 programmed, along with interface and address decoders in order to respond
to memory and input/output address ranges that are guaranteed to be
mutually exclusive to other devices forming part of the system. PCI
configuration is implemented using address space 0800H to 08FFH thereby
insuring that compatibility is retained with other environments of this type. PCI
15 bridges 416 and 417 also require the implementation of two hundred and fifty
six configuration registers, utilising two, thirty two bit registers located at
addresses 0CF8H and 0CFCH within the address space of processor 412.
These registers may be identified as the configuration address register and
the configuration data register.

20 The configuration registers are accessed by writing bus number,
physical device number, function number and register number to the address
register. Thereafter, an input/output read or write is performed to the data
register. The configuration address register only latches data when the host
processor 313 performs a full thirty two bit write to the register. Any eight or
25 sixteen bit access within this double word will be passed directly on to the
PCI bus as an eight or sixteen bit PCI input/output access.

Each bridge 417, 419 includes a set of configuration registers residing
with it's assigned range of two hundred and fifty six configuration locations to
permit tailoring of the bridge's functionality. The first sixty four configuration
30 registers are set aside for a predefined configuration header 701, including a
device identification, a vendor identification, a status register and a command

register. Bit one of the command register is set to enable memory access, such that the PCI bridge will respond to PCI memory accesses. An eight bit register 702 contains a number for the respective secondary PCI bus, assigned by the configuration instructions. A system re-set clears this register, whereafter reconfiguration by the configuration instructions is required in order to re-establish functionality.

Command/status register 703 provides for selection of operational characteristics. With bit two of this byte set, the bridge is unable to respond as memory on it's second bus. Memory base address 707 and memory limit address 705 are specified to define a range of memory addresses which, when generated on the primary bus 413, will result in a response being made by the respective PCI bridge. Thus, this range of addresses identifies a non-addressable range which allows the control processor to command instructions to the disc array 203. Similarly, memory accesses outside this specified range are ignored by the bridge, thereby providing the required isolation between the primary and secondary buses.

The PCI bridges are configured to allow processing unit 414 to issue command instructions to the network card 412, the application card 413 and the disc card 411 within a limited range of memory space. Consequently, the PCI bridges are not available for the transfer of image data between the secondary buses and the primary bus and a transfer of this type must take place via the router 402.

Color-space converter 403 is detailed in Figure 8 and includes a conventional digital converter matrix 801. The converter matrix 801 receives each input sample, multiplies samples by stored coefficients and then adds appropriate components in order to effect a color-space conversion. Thus, in typical applications, conversions are effected between YUV representations of color and similar RGB representations.

The conversion process is complicated by the fact that U and V color difference signals are often conveyed at a lower sampling rate and their associated Y component, while RGB samples are produced at a rate

compatible with said Y samples. Digital video signals having reduced bandwidth color components are often designated 4:2:2 to distinguish them from equally sampled components, represented as 4:4:4. The converter matrix 801 is configured to receive and produce samples in accordance with the 4:4:4 standard, therefore it is necessary to effect up-sampling or down-sampling of the color difference signals either on the input to the converter matrix or the output of the converter matrix, depending on the direction of conversion taking place. To avoid the need to duplicated converter circuitry, the color-space converter 403 is provided with a programable gate array, such as the 3K device manufactured by Xilinx of San Jose, California, USA.

The converter 403 includes a sample converter 803 arranged to up-sample U and V components to produce RGB samples or to down-sample RGB components to produce Y, U and V output samples. Y samples do not require down-conversion therefore the sample converter 803 also includes a delay device configured so as to maintain the Y samples in phase with down-sampled U and V components. An input from the router 402 is supplied to gate array 802 over an input bus 804. If the input samples are in RGB format, the gate array 802 is instructed, over a control channel 805, to direct said samples to the converter matrix 801. The converter matrix 801 converts the RGB samples to YUV samples which are in turn returned to the gate array 802 via bus 806. Upon receiving these samples over bus 806, the gate array directs said samples to the sample converter 803 which reduces the rate of the U and V samples to produce samples falling within the accepted 4:2:2 protocol on an input bus 807. The gate array receives input samples on bus 807 and directs these to an output bus 808 which is in turn directed to the router 402.

Alternatively, the color-space converter 403 may be configured to convert YUV samples to RGB samples. The gate array 802 is instructed, via control channel 805, to the effect that it will be receiving YUV samples. The incoming YUV samples on bus 804 are firstly directed to the sample converter 803 which up-samples the U and V components to produce 4:4:4

YUV samples which are returned to the gate array 802 on input bus 807. Within the gate array 802, said input samples are directed to the converter matrix 801, arranged to generate RGB representations of these samples which are in turn returned to the gate array 802 by input bus 806. Within gate
5 array 802, the samples received on bus 806 are directed to the output bus 808. Thus, it can be appreciated that the gate array 802 allows samples to be directed to both the converter matrix 801 and the sample converter 803 in either order.

Proxy generator 404 is detailed in Figure 9. Data is supplied from the
10 router 402 to the proxy generator 404 over thirty-two bit bus 901, consisting of eight bits allocated for the red component, eight bits allocated for the green component, eight bits allocated to the blue component and eight bits allocated to a control channel also known as a keying channel or an alpha channel. Bandwidth reduction of the control channel does not have meaning,
15 therefore the eight bit red, green and blue components are supplied to respective pre-filters 902, 903, 904 and the control bus is effectively terminated at 905.

The pre-filters provide bandwidth reduction for relatively large images, such as those derived from cinematographic film. When broadcast video
20 signals are received, no pre-filtering is necessary and bandwidth reduction is performed exclusively by a re-sizing device 905 which, in the preferred embodiment, is a GM 833X3 acuity re-sizing engine manufactured by Genesis Microchip Inc. of Ontario Canada.

The re-sizing device 905 receives data over lines 906 and 907
25 specifying source image size and target image size respectively. Outputs from pre-filters 902, 903 and 904 are supplied to respective buffering devices 908, 909 and 910. Each buffering device includes a pair of synchronised field buffers, such that a first field buffer 911 is arranged to receive a field of data from pre-filter 902 while a second field buffer 912 supplies the previous field
30 to the bandwidth reduction device 905.

Bandwidth reduction device 905 receives outputs from each of the

buffering devices 908, 909, 910 and effects bandwidth reduction upon each of the red, green and blue components in response to the programmed reduction size. In this way, the bandwidth reduction device has access to the data stored in one of the field buffers, representing the source buffer
5 throughout a field period. Similarly, throughout this period output values for red, green and blue components are supplied to respective output buffering devices 913, 914 and 915. Again, each output buffering device includes a pair of co-operating field buffers 916 and 917.

The outputs from the buffering devices 913, 914 and 915 are
10 reassembled into a thirty-two bit output bus 916, with its eight bit control bytes effectively set to nil.

The re-formatters 406 and 407 are implemented primarily using logic cell arrays, such as the Xilinx XE3000. The devices are field programmable gate arrays configured to replace conventional TTL Logic's devices and
15 similar devices which integrate complete subsystems into a single integrated package. In this way, a plurality of packing and unpacking configurations may be programmed within the device which are then selectable, in response to commands issued by the control processor 412, for a particular packing or unpacking application.

20 User logic functions and interconnections are determined by configuration program data stored in internal static memory cells. This program data is itself loaded in any of several available modes, thereby accommodating various system requirements. Thus, programs required to drive the devices may permanently reside in ROM, on an application circuit
25 board or on a disk drive. On chip initialization logic provides for automatic loading of program data at power-up. Alternatively, the circuit may be associated with an XC17XX chip available from the same source, to provide serial configuration storage in a one-time programmable package.

30 Within the device, block logic functions are implemented by programmed look-up tables and functional options are implemented by program controlled multiplexes. Interconnecting networks between blocks

are implemented with metal segments joined by program controlled pass transistors.

Functions are established via a configuration program which is loaded into an internal distributed array of configuration memory cells. The configuration program is loaded into the device at power-up and may be reloaded on command. The logic cell array includes logic and control signals to implement automatic or passive configuration and program data may be either bit serial or byte parallel.

The static memory cell used for the configuration memory and the logic cell array provides high reliability and noise immunity. The integrity of the device configuration is assured even under adverse condition. Static memory provides a good combination of high density, high performance, high reliability and comprehensive testability. The basic memory cell consists of two CMOS inverters plus a pass transistor used for writing and reading cell data. The cell is only written during configuration and only read during read-back. During normal operation the cell provides continuous control and the pass transistor is off and does not affect cell stability. This is quite different from the operation of conventional memory devices, in which the cells are frequently read and rewritten.

An array of configurable logic blocks provide the functional elements from which the packing and unpacking logic is constructed. The logic blocks are arranged in a matrix so that 64 blocks are arranged in 8 rows and 8 columns. Development software available from the manufacturer facilitates a compilation of configuration data which is then downloaded to the internal configuration memory to define the operation and interconnection of each block. Thus, user definition of the configurable logic blocks and their interconnecting networks may be done by automatic translation from a schematic logic diagram or optionally by installing a library of user callable macros.

Each configurable logic block has a combinational logic section, two bistables and an internal control section. There are five logic inputs, a

common clock input, an asynchronus reset input and an enable clock. All of these may be driven from the interconnect resources adjacent to the blocks and each configurable logic block also has two outputs which may drive interconnected networks.

- 5 Data input from either bistable within a logic block is supplied from function outputs of the combinational logic or from a block input. Both bistables in each logic block share asynchronus inputs which, when enabled and high, are dominant over clocked inputs.

- 10 The combinational logic portion of the logic block uses a 32 by 1 bit lookup table to implement Boolean functions. Variables selected from the five logic inputs and two internal clock bistables are used as table address inputs. The combinational propagation delay through the network is independent of the logic function generated and is spike free for single input variable changes. This technique can generate two independent logic
15 functions of up to four variables.

- Programmable interconnection resources in the logic cell array provide routing paths to connect inputs and outputs into logic networks. Interconnections between blocks are composed of a two layer grid of metal segments. Pass transistors, each controlled by a configuration bit, form
20 programmable interconnection points and switching matrix's used to implement the necessary connections between selected metal segments and block pins.

- The re-programable nature of the device as used within the reformatting circuit 406 results in the actual functionality of these devices
25 being re-configurable in response to down-loaded instructions. The devices essentially consist of many registers and as such provide an environment in which the reformatting links may be effectively "hard-wired" in preference to being assembled from the plurality of multiplexing devices.

- An example of the functionality within reformatting device 406 is
30 illustrated in Figure 10. Input signals from router 402 are supplied to a width selector 1001, arranged to separate RGB sub-words into eight bit